

AMD EPYC™ 7002 SERIES PROCESSORS WEATHER MODELING WITH WRF



AUGUST 2019

AMD EPYC for HPC

Utilizing the x86-architecture, and built on 7nm technology, the AMD EPYC™ 7002 Series processors bring together high core counts, large memory capacity, extreme memory bandwidth and massive I/O with the right ratios to enable exceptional HPC workload performance.

Standards Based Architecture

Continuing the AMD commitment to industry standards, AMD EPYC™ 7002 Series processors offer you a choice in x86 architecture. x86 compatibility means you can run most x86 based applications on AMD EPYC processors.

Exceptional Scalability

Scaling is critical to HPC applications. AMD EPYC 7002 Series processors provide high bandwidth between nodes with support for PCIe Gen 4 enabled network devices. Within node, take advantage of up to 64 cores per socket, including 8 memory channels utilizing speeds up to DDR4-3200². Add incredible floating point and integer compute within each core and the AMD EPYC 7002 series delivers exceptional performance and scalability for HPC.

Fully Tested and Validated

AMD's broad partner ecosystem and collaborative engineering provide tested and validated solutions that help lower your risk and total cost of ownership.

Weather Research & Forecasting (WRF)

The Weather Research and Forecasting (WRF) Model is a next-generation mesoscale numerical weather prediction system designed for both atmospheric research and operational forecasting applications.

AMD EPYC™ 7002 Processors: Architectural Innovations Deliver Exceptional Performance and Scalability

The high-performance computing (HPC) market has grown to a point where it is a critical component of new technology advancements in academia and a wide array of industries in both the public and private sectors. Scientific research, public health, climate modeling, as well as oil and gas exploration are just a few examples where HPC is the driving force behind new innovations and knowledge discovery.

7 nm	PCIe® Gen 4	DDR4 3200
64 Cores per socket	128 PCIe® Gen 4 lanes per socket	8 Memory channels per socket
World's first 7 nm x86 server CPU Highest available core count ¹ to maximize parallelism	World's first PCIe® Gen 4 ready x86 server CPU ² Doubles the bandwidth of the previous generation	World's first x86 architecture with DDR4 3200 ² Up to 4 TB of memory capacity per socket

The second generation of the AMD EPYC™ processor extends AMD innovation leadership for HPC. Built with leading-edge 7nm technology, the AMD EPYC™ SoC offers a consistent set of features across a range of choices from 8 to 64 cores, including 128 lanes of PCIe® Gen 4² and 8 memory channels with access to up to 4 TB of high-speed memory.

The AMD EPYC™ 7002 Series processor's innovative architecture translates to tremendous performance and scalability for HPC applications, offering you a choice in x86 architecture while optimizing total cost of ownership.

Weather Modeling with WRF

The Weather Research and Forecasting (WRF) Model is a next-generation mesoscale numerical weather prediction system designed for both atmospheric research and operational forecasting applications. It features two dynamical cores, a data assimilation system, and a software architecture supporting parallel

computation and system extensibility. The model serves a wide range of meteorological applications across scales from tens of meters to thousands of kilometers.

For researchers, WRF can produce simulations based on actual atmospheric conditions (i.e., from observations and analyses) or idealized conditions. WRF offers operational forecasting a flexible and computationally-efficient platform, while reflecting recent advances in physics, numerics,

and data assimilation contributed by developers from the expansive research community. WRF is currently in operational use at National Centers for Environmental Prediction (NCEP) and other national meteorological centers as well as in real-time forecasting configurations at laboratories, universities, and private companies. Memory bandwidth is a critical factor in maximizing performance of weather modeling workloads. AMD EPYC server processors' exceptional memory bandwidth helps ensure that you get the most out of your system, optimizing execution time and overall utilization of your deployment.

As workloads demand performance requiring more processor cores, the communications between processor cores becomes critical to efficiently solving the complex problems faced by customers. As cluster sizes increase, the communication requirements between nodes rises quickly and can limit scaling at large node counts.

AMD addresses this scaling limitation by partnering with leading network providers, such as Mellanox®, to offer high performance solutions based on PCIe Gen 4. PCIe Gen 4 allows incredible performance on Infiniband HDR 200Gb/s fabric when run on 2nd Generation EPYC processor-based clusters.

Performance Benchmarks and Testing

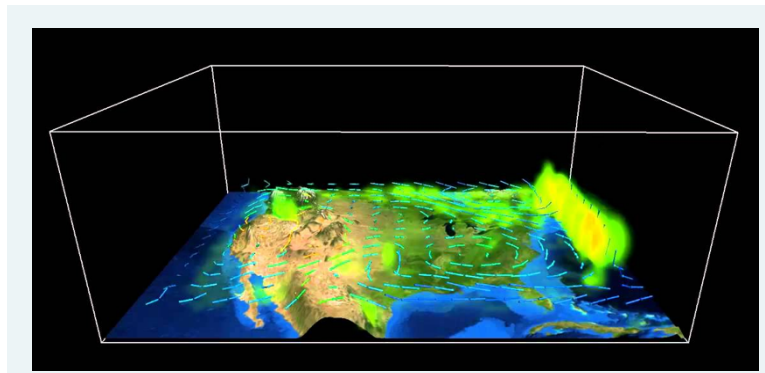
This document focuses on performance and scaling of the EPYC 7002 Series Processors. Testing was performed on a cluster of dual-socket EPYC 7742-based systems and dual-socket EPYC 7542-based systems. Each EPYC™ 7742 processor has 64 cores with a base frequency of 2.25 GHz and a boost frequency of 3.4 GHz. Each 7542™ processor has 32 cores with a base frequency of 2.9 GHz and boost of 3.4 GHz.

The compute nodes in the cluster are each populated with 1 DIMM per channel of 64-GB, dual-rank, DDR4-3200 DIMMs from Micron®, for a total of 1TB of memory per node.

A Mellanox® ConnectX-6 200 Gb/s HDR InfiniBand adapter, utilizing EPYC processors' support for PCIe Gen 4, is also populated on each EPYC processor based system.

Single-node testing was performed across all platform configurations and multi-node scaling was tested on the EPYC 7542 processor.

WRF benchmarks provide a basis of evaluating hardware performance. Standard models are provided that represent typical usage. The benchmarks used were conus12km and conus2.5km. All runs were in hybrid MPI/OpenMP with 4 OMP threads per rank.



The EPYC Advantage: AMD EPYC server processors offer 8 memory channels of DDR4-3200² and support for up to 4 TB of memory per processor, yielding exceptional memory bandwidth and capacity.

Tested Hardware and Software Configuration

AMD 2 nd Generation EPYC Compute Nodes		
CPU	2 x EPYC 7742	2 x EPYC 7542
Cores	64 cores per socket (128 per node)	32 cores per socket (64 per node)
Memory	Micron 1 TB (16x) Dual-Rank DDR4-3200, 1DPC	
Network Adapter	Mellanox ConnectX-6 200Gb/s HDR InfiniBand x16 PCIe [®] Gen 4	
Storage: OS Data	1 x Micron 1100 256 GB SATA 1 x 1 TB NVMe	
Software		
OS	RHEL 7.6 (3.10.0-862.el7.x86_64)	
Mellanox OFED Driver	MLNX_OFED_LINUX-4.5-1.0.1.0 (OFED-4.5-1.0.1)	
Network		
Switch	Mellanox 200Gb/s HDR InfiniBand Switch (MQM8790)	
Configuration Options		
BIOS Setting	NPS = NPS4, SMT = Off, Boost = On, X2APIC = On, Determinism Slider = Performance, Preferred IO=Enabled	
OS Settings	Governor=Performance, CC6 = Disabled	

AMD 1 st Generation EPYC Compute Nodes		
CPU	2 x EPYC 7601	
Cores	32 cores per socket (64 per node)	
Memory	256GB (16x 16GB Dual-Rank) DDR4-2666	
NIC	Mellanox ConnectX-5 EDR 100Gb InfiniBand x16 PCIe	
Storage: OS Data	1 x 256 GB NVMe 1 x 1 TB NVMe	
Software		
OS	RHEL 7.6 (3.10.0-862.el7.x86_64)	
Mellanox OFED Driver	MLNX_OFED_LINUX-4.5-1.0.1.0 (OFED-4.5-1.0.1)	
Network		
Switch	Mellanox HDR 200Gb/s Unmanaged Switch (MQM8790)	
Configuration Options		
BIOS Setting	SMT = Off, Boost = On, Determinism Slider = Performance, Global C-State Control = Enabled	
OS Settings	Governor=Performance, CC6 = Disabled	



WRF Compilation

WRF version 3.8.1 was compiled from source on RHEL 7.6 using Intel Parallel Studio 2019 compiler. WRF has dependencies on a few libraries. The following were used. OpenMPI 3.1.3, netcdf-4.1.3, zlib-1.2.7, libpng-1.2.50, jasper-1.900.1. All libraries were also compiled from source with Intel Parallel Studio 2019 compiler. The default optimization flags were used. No further compile time optimizations were done.

WRF Performance: Single Node Performance

Single node performance was compared between the 1st Generation EPYC 7601 processor (32 cores) and the 2nd Generation EPYC 7542 (32 cores) and EPYC 7742 (64 cores) processors. The industry standard WRF CONUS 12km benchmark was run on each configuration.

CONUS12km is a 48-hour, 12km resolution case over the Continental U.S. (CONUS) domain October 24, 2001 with a time step of 72 seconds. The benchmark period is hours 25-27 (3 hours), starting from a restart file from the end of hour 24 (provided).

Figure 1 shows outstanding performance on the CONUS12km model. The 32-core EPYC 7542 shows a ~19% generational improvement, while the 64-core EPYC 7742 shows an impressive ~65% improvement. The next section will show scaling of the EPYC 7542 on the larger CONUS2.5km model.




Figure 1

WRF: Multi-node Scaling

Conus2.5km is the larger of the two benchmarks. It is the latter 3 hours of a 9-hour, 2.5km resolution case covering the Continental U.S. (CONUS) domain June 4, 2005 with a 15 second time step. The benchmark period is hours 6-9 (3 hours), starting from a restart file from the end of the initial 6-hour period.

Conus2.5km model was run in a clustered environment, still using a hybrid MPI/OpenMP configuration, with four OMP threads per rank. The scaling of the 32-core EPYC 7542 processor continues to maintain exceptional efficiency through at least 8 nodes (512 cores), as shown in Figure 2.

Exceptional scaling of this magnitude can be attributed to cache effects fitting more of the model into the processors' cache. For some workloads, this means that instructions that might have otherwise been delayed waiting for a memory request to complete, can complete much faster with data already sitting in the cache, effectively improving IPC (Instructions Per Clock).

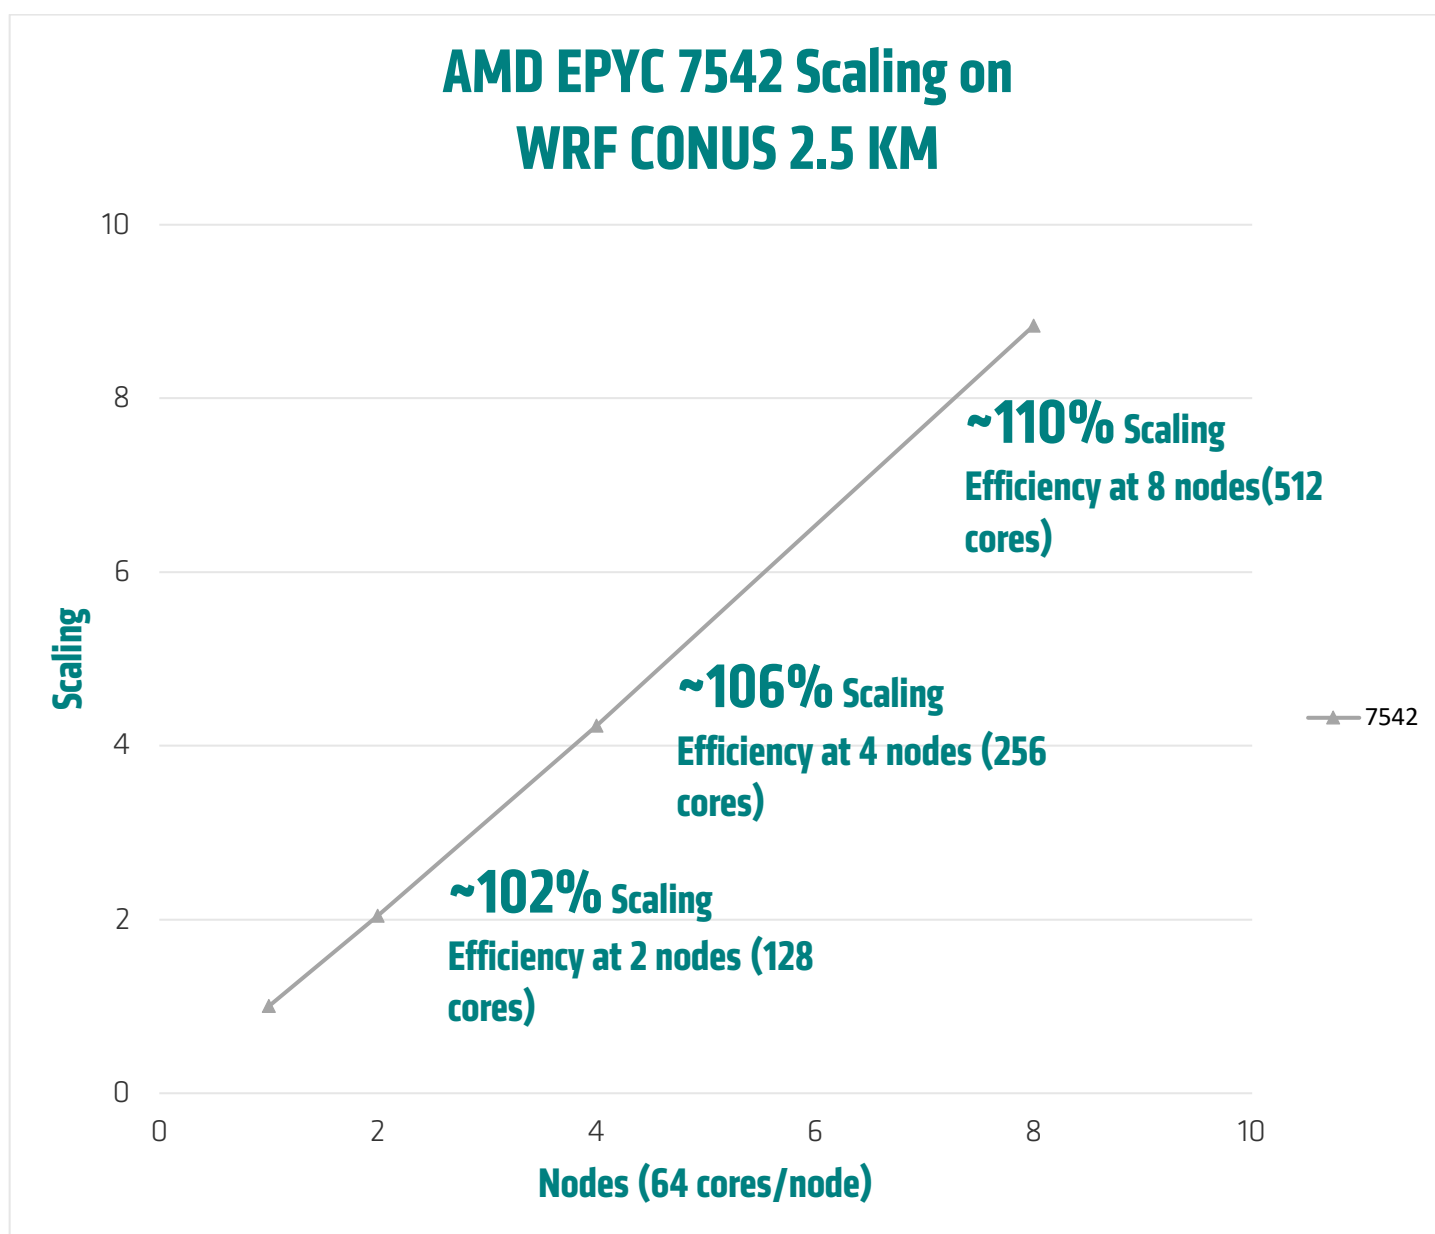


Figure 2

Summary

WRF benchmarks were conducted on single-node, two-socket systems of 2nd Generation AMD EPYC 7742 and 7752 processors and 1st Generation AMD EPYC 7601 processors. The smaller CONUS 12KM model was run on single-nodes and then CONUS 2.5KM was scaled out on the 2nd Gen EPYC 7542 processors using multiple systems. Planned future testing will explore the exceptional scaling the 2nd Gen EPYC processors can achieve. WRF scales well across the benchmarks to many cores.

Conclusion

Scale-out testing on the EPYC cluster shows impressive results on these benchmarks. Pure performance was highest with the 64-core 2nd Gen EPYC 7742. Per-core performance was highest with the 32-core 2nd Gen EPYC 7542. Whether you need the dominating system level performance and density of the EPYC 7742 or the equally dominating per-core performance of the EPYC 7542, all products offer exceptional core IPC and memory bandwidth, and both provide your organization a significant advantage. Customers can pick the optimal part based on their unique requirements.

AMD empowers the development of fast, accurate weather modeling and simulations running on cost-effective clustered systems.

For more information about AMD's EPYC line of processors visit: <http://www.amd.com/epyc>

For more information about WRF visit: <https://www.mmm.ucar.edu/weather-research-and-forecasting-model>

FOOTNOTES

1. Best-in-class based on industry-standard pin-based (LGA) X86 processors. NAP-166.
2. Some supported features and functionality of second-generation AMD EPYC™ processors (codenamed "Rome") require a BIOS update from your server manufacturer when used with a motherboard designed for the first-generation AMD EPYC 7000 series processor. A motherboard designed for "Rome" processors is required to enable all available functionality. ROM-06.

DISCLAIMER

The information contained herein is for informational purposes only and is subject to change without notice. While every precaution has been taken in the preparation of this document, it may contain technical inaccuracies, omissions and typographical errors, and AMD is under no obligation to update or otherwise correct this information. Advanced Micro Devices, Inc. makes no representations or warranties with respect to the accuracy or completeness of the contents of this document, and assumes no liability of any kind, including the implied warranties of noninfringement, merchantability or fitness for particular purposes, with respect to the operation or use of AMD hardware, software or other products described herein. No license, including implied or arising by estoppel, to any intellectual property rights is granted by this document. Terms and limitations applicable to the purchase or use of AMD's products are as set forth in a signed agreement between the parties or in AMD's Standard Terms and Conditions of Sale. GD-18

©2019 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, EPYC, and combinations thereof are trademarks of Advanced Micro Devices, Inc. PCIe and PCI Express are registered trademarks of PCI-SIG Corporation. Micron, the Micron logo, and the Micron symbol are trademarks of Micron Technology, Inc. Mellanox and the Mellanox Technologies logo are registered trademarks of Mellanox Technologies. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.